

ASI Technology Summit Q1 2021
Micron Memory Solutions for Today, Tomorrow, and Beyond

0:00

Starting, all attendees are in listen only mode.

0:03

Hi. Good morning, or good afternoon, everyone. This is Kent Tibbils with ASI. I want to welcome you all to day three of our ASI technology summit.

0:14

As you can see today, we have a nice big group with us from Micron.

0:20

So, Micron is going to be talking on a variety of topics including doing an introduction of some new team members that are going to be aligned for supporting ASI.

0:32

So, before we get started with the Micron team and their presentation, there's a couple of things that I wanted to go over.

0:40

First, following up from yesterday, we had a great presentation from ASUS. They talked about their family of motherboards, and how they categorize their products. And they saw we saw some new technologies that I don't think a lot of us were aware, that ASUS had some things that make their boards unique.

0:59

If you didn't get a chance to see that, we've provided the recorded webinar for you guys to review along with the slide deck which I just got this morning and I'll be sending out to you guys later today, and that was for yesterday's presentation. So the other thing from yesterday is we had some prizes that we gave away and we said, we're going to announce the winners of those prizes this morning.

1:22

So I'm going to go ahead and do that right now just to let everybody know. So yesterday, ASUS was giving away a 32 inch touch.

1:32

Monitor an LED monitor, and the winner of that is going to be sent to Anthony Wu.

1:38

Anthony is one of our customers up north in Canada, so congratulations, Anthony will be connecting by e-mail to make sure that we can get you that monitor.

1:49

Also, in addition to that, again, because we had such great turnout yesterday, and a great participation from you guys, with questions and everything, we're going to give away some additional \$50 gift cards, so we've got 10 of them to give away, and I'm going to read off the list of those winners here really quick.

2:09

We have David Bacon, Allison, both Schreyer, Geoff Brossard, Mike, ..., Shawn Herrera, Karl Hold, May, Liew, William Richmond, Gabriel, Trap, Jani, and Don Young. Congratulations all of you lucky winners of that \$50 card.

2:34

I'll be sending you some e-mails after everything is done.

2:38

So, we'll get all the T's and C's crossed, and everything, so we can get those sent out to you guys.

2:44

So, before we start, one thing I want to let you guys know is for questions. If you have questions, you can click on the question box and go ahead and type in the question and I'll be able to ask those to our group.

2:57

As we go through the presentation, we have things kind of broken up into segments based on gaming and DRAM and some market information. So, at the end of each presentation, we're going to take a pause for a minute to see if we have any questions for that specific presenter.

3:16

So, if you guys do have questions, so if you have questions for, you know, Jay, or you have questions for Ryan, after they've done their presentation, make sure you type that question in while they're presenting, so that we can go ahead and ask at the end of their session, and then we'll move on to the next one.

3:33

So if you have questions, type them in, and we'll make sure we'll get those answered for you, OK.

3:38

So, with that said, you know, nobody wants to hear me ramble on forever really came here to hear from Micron and the whole group.

3:46

That's here.

3:47

So, I'm going to go ahead and introduce Paul, and let Paul go ahead and kick everything off.

3:54

For team, Micron, Paul, take it away, the floor is yours.

3:58

Perfect. Thanks, Kent. Appreciate it.

4:00

As Kent mentioned I'm Paul Bedke, I'm distribution sales manager for Micron and I wanted to thank everyone for jumping on. I know we're all busy, and we have a lot of zoom interaction.

4:10

We appreciate your time and appreciate ASI putting this on their great partner of ours and appreciate their efforts. As Kent mentioned, have a lot of info to cover. And I've got some experts on here.

4:23

So I don't want to get too much into their time, but probably my biggest announcements and information going forward is I've got a new team member two weeks in.

4:31

I believe, David Lefstein, two weeks with Micron, but several years of experience in the distribution and storage world. So super excited to have him on board and I'm going to turn it over to him.

4:44

Thanks, call, Appreciate it. Very excited to be back with a partner like ASI and work with Micron.

4:50

And as Paul mentioned, we're excited to introduce our experts here, Andrew, Ryan and Jake. I think we have some exciting things to share with you guys.

4:58

And, without further ado, I'm going to push it over to Andrew and start going through the process.

5:03

Thinkable.

5:04

All right. Welcome, everybody. So we're gonna start out talking about the market and what's happening and if you've been a part of the memory of the storage industry, really, at any point over the last 20 or 30 years, you know that we go through ups and downs. It's very cyclical in that manner.

5:20

And so on, an update on what's happening today, and what the expects is going to happen over the next couple.

5:28

So, on this first slide, here, it's really just a collection of a lot of what the analysts are talking about, and when we take a look, Tom.

5:39

We will look at it in both a memory and storage site, or a DRAM, and, and, and side. And on the DRAM side, what you're going to see is you're gonna see tighter supply and typically a higher prices as we go throughout the year. And I remember, even though there are a lot of modular ... module manufacturers, there's really only three D ram suppliers so there. There's not a lot in this that are servicing this industry. And so, typically, what happens, especially when we, when we come out of a year where we've had ups and downs of supply and demand due to co-lead and lots of other things, what we're seeing is the uptick in terms of.

6:17

demand versus supply. And so on the D ram side, you're going to start seeing the, you know, ASPs going out throughout the year.

6:25

On the ..., on the storage side, you're going to see something similar. But, maybe to a lesser extent. So, NAND, typically, if you go back over the last, you know, anywhere from 3 to 10 years, every year, we see a pretty sharp decline in ASPs, usually in the realm of about 20% per year of negative ASP. When we go through times the constrain, what you'll see is, is you may see some increase, but typically it's more of a leveling out. And so where we're at right now is that we're going to see a little bit of a price increase.

6:57

It's certainly not going to be on the level of ram, And really that was a lot of the same factors on DRAM. But we do have one thing that the industry is running into. It's not really affected micron at this point. But we have had some control slash controller companies that have run into some issues during the big snowstorm and low temperature period. In Texas, Samsung's fab for producing flash controllers, went down for an extended period of time and even though, you know, there are obviously a competitor of ours. There's only so much supply in the industry and when somebody locked, like Samsung has any sort of pickup. You're gonna see that propagate to everybody else. So, in terms of the Outlook going forward, make sure that if you have any sort of opportunities that you're able to and, you know, as best you can, communicate with us as quickly as possible so that we can help, you know.

7:49

You'll be able to address the sort of demand that you and your customers are generating.

7:55

So, as we move on to the next slide, um, 2021 in terms of the mark in terms of more of the technology, there's going to be a big shift going on and you're gonna see it on the memory and the storage side. On the memory side, everybody is ramping up and getting ready for DDR five and. and we've got two experts that are going to follow me that if you have any questions on that, you can certainly certainly ask them. But it's a big transition now, when you talk about any sort of system. Whether it's a gaming system or a server within the data center, you have to have balance within that system. And so you can't just make one component or one sub system faster and expects benefits across the board. You really has to be able to to grow all of the different pieces, whether it's compute or networking or in the case of of what we're talking about today, memory and storage. And as we move into PCIE gen foreign DDR five, what you're going to see in the table on the right is roughly a doubling of performance. And so that's a big big jump.

8:58

On the PCIE, Gen four side. You may say, well, we've had Gen four out for a little while, actually, more than a year now. But it's been really limited to AMD systems. As we go throughout 2021, you're gonna see Intel get on board. They're obviously a very, very large part of the Market, whether it's in consumer data center, gaming, any of those pieces and intelligent, very big player. and so, with them getting on board with Gen four. And as you see, DDR five platforms become available throughout the end of the year. There's just a lot of transition going on the industry right now. And so, we wanna make sure that you're prepared for those shifts, so that you're able to support your customers.

9:37

So it's not just the industry standards that we're talking about. We want to make sure that you understand some of the things that ... working on. And so, on the storage side, on the NAND flash side, we announced late last year, that we were the first supplier, to hit 176 layer NAND,

And so you may not be a semiconductor unmanned expert, but anytime you hear somebody like Micron, or Intel, or Samsung talk about numbers of layers with, Dan.

10:06

What it means is that we're able to put more bits on a wafer. And the more bits we pull them away for those, the higher the availability of better supply. We get, while at the same time, you're reducing cost. And so being the first to 176 layer may not mean anything to you right now, but as you look at our storage portfolio going forward, over the next couple months, to the next 12 months, you're gonna see this technology. Make it into our SSDS and you're gonna be able to get them at a lower cost. And really, this is the best of all worlds. It's, it's not only lower cost. But for this specific announcement, we're also, we also had a really large improvement in performance in terms of latency.

10:47

And so, you know, it's something, we were the first to market, and we weren't the first to market by a few weeks or a few months. We're actually, in some cases, 3 to 1.25 ahead of our competition, so that puts us in a really good spot over over the next year. So, something we're very proud of that you'll start seeing you know the fruits of that very soon.

11:10

And then, finally, before we jump into the DRAM and the memory side, we definitely wanted to make sure that we covered just a little bit on the SSD side.

11:21

So, let's take a look at this final slide here.

11:23

And what we'll see is we do have a broad portfolio of data center SSDS, the ones that were shown here whether you want something that is incredibly high performance low latency really for the the most stringent you know the most demanding applications with products like our 9300 which is in NVME base SST, That goes all the way up to 15 terabytes. We also have products like the 7300, which is more mainstream, which is the balance of price, performance, and features. And then, if you still have customers that are on Sarah, will be supporting Sada for a long time, going forward. So, whether it's a mainstream 5300, are q.l. sea bass 50 to 10, which is really meant as a hard drive replacement. It's going to be the lowest cost SSD in our portfolio. So if you have any SSE opportunities please please make sure to reach out to myself and David. Paul, any of us are more than happy to sit down with you or your customer to discuss those opportunities.

12:20

But going forward, the rest of this time we're going to be. We're going to be talking about, you know, our DRAM offerings, and if there are any questions, more than, more than happy to answer them right now.

12:30

And if not, we'll, we'll give it away to Brian and Jake actually got a real quick question Here for you. Regarding the 186 layer ..., is that SLC MLC both, the, what was the NAND on that?

12:48

It's going to be, the initial products will be TLC based.

12:52

So TLC base. Great.

12:54

Yes, OK, good question.

12:59

Let's see, are there any power consumption savings to be had with these new technology? So as the layers increase and ..., are there better power savings that are occurring also?

13:13

Yes. Typically, as we move from generation to generation, you will see some you will see power savings in a in a in a performance performance comparison, but you what we're also battling right here is as we go from things like PCIE Gen three to Jen for the products get faster. And typically on an SSD, the faster you go, the more power you draw. And so it is a give and take, So if you had a product that was is, was equal in performance, you would see a decrease in the power draw. But as we go to Gen three, and we doubled, the speeds were not doubling the power. But you will see a power N So, it's very much a classic example of you get what you pay for, Right? And if you want that higher speed, you're typically going to draw more power.

14:02

Quick question.

14:05

I think we're good.

14:06

So I got one right here that maybe, hopefully you can answer this one, But on the 55210, is that a ...? So there's, no, is there DRAM cash on that one?

14:20

So so in order to get the cost down, we went with Q L C, nan, so that's that's how we get the lower cost profile. But it is a sole DRAM enabled solution. So the difference between the 62 10 and 5300 is really just demand. That is the major difference. Architecturally, it's the same slash controller a six very similar firmware code base and the same DRAM products.

14:46

Yeah, that's actually a great question. So, as we get ready to transition here, so this is my way of allowing the next presenter to kinda get ready.

14:56

There was some article posted that a lot of you guys are gonna be going to SSDI that don't have cash as a way to keep the cost down on our systems and on their notebooks. So, I think for you guys out in the channel, that as you're looking forward, you're going head to head against these brands.

15:14

It's a great question and something to really look at from the good guys and make sure that as their SSP is to have cash, because it'll be a Performance Advantage that you can have in your system verse verse.

15:27

There's, as they're trying to maintain, cost points that they need to achieve, the ad group, and you will typically see that on clients and consumer. When you get into the value segments, there, it's,

it's, it's going to be very dominated by DRM solutions because for every every terabyte SSD storage, you typically have a gigabyte of ram. And so those costs add up especially in higher capacities. Andrew, you're going to be with us all the way to the Right.

15:59

So if we have more questions for you, OK, great, so Absolutely, go ahead, guys, Keep sending in your questions.

16:04

We're gonna go ahead and move to our, to our next presenter and I believe if I have my notes corrected about Ryan.

16:13

So, Brian, you're up.

16:17

That's correct, yeah, hopefully everybody can hear me, microphone check?

16:20

We can hear you great.

16:22

Perfect, alright. Yeah, absolutely, let's start. Let's talk about the fun stuff.

16:27

So our are gaming ram and performance ram. So we have two product lines will talk about both product lines, Ballistics, Ballistics Max.

16:36

We'll talk about some of the There records, we've set the performance around it, and then we'll get into actually how we're doing this and what that differentiates us from the competition. It's actually some, some pretty interesting stuff.

16:47

So first off, let's talk about the crucial ballistics line. So this really is our main stream line. This is a, this is, I'd say, is covers about 90% of, of our director in the gaming markets that goes out there, we've set this product line up to really hit the majority of the market. So when you look at speeds, he's got speeds from 2666 all the way up to 3600 Seals 16. That speeds very important for us today, right, because with AMD being more and more popular, 3600 C L 16, really, is a sweet spot.

17:21

We wanted to make sure we had the majority of our market there are captured within this, this, this value segment of our gear.

17:32

So this is offered in 8, 16, and 32 gigabyte modules. We do offer RGB solutions on it, in, in all three colors, starting with 3200 speed.

17:44

So RGB today, obviously is a big deal in gaming. I've got it behind me, Mom Intuit, for sure. You know, gaming is becoming a very individualized. People want to make a system of

their own. And RGB is important. So, we do offer that and it is compatible with DOD and software out there.

18:03

So I heard that a suit was, was on deck yesterday. Aces or we'll work with this MSAs, mystic like and so on. We also have our own utility. If he's got an older chipset you may want to run the RGB on this. It is a blasted aluminum heat spreader.

18:20

Again, more towards esthetics in this product, but we do have black, white and red in both RGB and non RGB.

18:27

So if you don't want the RGB, you don't have to have it which which is, you know, some people prefer that.

18:37

Let's go to the next slide.

18:42

All right, let's get into ballistics Max first. I want to touch on some of the accomplishments we've had with this product.

18:48

So ballistics. Max, really is our high-end product gets our halo product, our performance product. Currently, today, and I believe, we've held this record, for about 6 or 7 months.

18:59

We, are the fastest memory on the planet at 6666 MCS with liquid nitrogen.

19:07

So, the ... have taken this product, and currently today we're, we're number one. I want to say we've broken about five records in the last year and a half or so, and we'll continue to drive those records. It's important to us because we learn about the product.

19:22

We can take that engineering and drive it back into the into the product line itself.

19:26

So, the consumer benefits two micron die. So, we are the only vertically Integrated Gaming DRAM Company, and I'll save that, that statement for the very end, and we'll talk about the benefits there.

19:41

Also, we are currently the fastest available DRAM on the market. Meaning that at an ETL level, you can go today, and, and, or we had available, are 50, 100 speed. That is the fastest on the market today. Obviously, that is a very much a halo product. It you have to have the right hardware to try to get that to work, But it's something that we wanted to put out there.

20:04

Mark?

20:09

Next slide.

20:11

Ballistics Max as a product line itself, we already touched on on the world records. Let's talk about speed. So this is offered from 4000 Holloway and all the way up to 5100 today, so 4040, 450, 100. 8 is 16. Gigabyte modules only on this product, and black only as well.

20:31

So being halo product we wanted to keep. We wanted to keep it fairly low skew count and really focus on the performance itself. There are a couple of key features on this product. There are quite different from anything that's out there.

20:45

In the upper left, you can see that three-d. printed light bar. This is kind of a neat feature that we carried over from an older product, where we offer the base file for this product.

20:55

You can three-d. print, whatever you want on top, slide in the slide, it in that top bar, and it will light up. So if you want your gamer tag or if you've got a system integrator, we've seen that before.

21:05

They want to say, put their integration name across the top of the memory and light it up. That's an option for this product. So something that's not necessarily part of performance, but an added benefit you get when you jump into this higher end product.

21:17

Then the other thing that's interesting here is thermal monitoring onboard. So, when you're pushing this stuff to the limit, obviously, it's gonna get hot, you're using up that voltage. We've got thermal monitoring onboard for this.

21:28

So, you can look at each individual dem and see what your temperatures are.

21:34

Then, speaking of thermals, the heat sink on this is quite robust. Unlike ballistics were ballistics Product Line where that's mostly esthetics are actually utilizing the heat spreader here, to kind of work out some of the heat. So, there is a thermal interface material between it and the heat spreader, This heat spreaders, a very large extrusion, so it's a very thick heavy in the hand.

21:56

The top bar on the non RGB is extruded aluminum. So, it really is truly there to wick some of that heat out.

22:04

Then, some other points I wanted to touch on this product, because lot of times, we just talk about gaming, and we kinda gets stuck in that red of gaming. This product has actually been utilized in a lot of interesting areas, as well. So, content creation of content creation is a big deal. Those that are running thread refers to do content creation. This product is perfect for that.

22:24

We've also seen areas where people have inquired about stock trading applications for this product, as well as VR. Obviously, that's getting very popular in automotive design and storefront design. We've actually seen this product utilized in that application as well. So, it's

important to keep an open mind when we, when we think about gaming memory, it's not just, you know, stuck in that box gaming, this stuff is, is utilized in a lot of other areas.

22:52

All right, Next slide.

22:57

OK, so, we talked about world records, we, you know, we talked about the fastest in the world, available on the market, it's, we're, we're able to do that for a couple of reasons, and it's definitely something our competitors can't do or they're not doing today.

23:11

So, as I said before, you know, we're the only vertically integrated memory company in gaming.

23:18

So, what are we doing here with this stuff? So, the first thing that we're doing is, we're going out of the at the die level.

23:25

We're bidding it, basically the wafer level, right? So we find candidates for ballistics. We segregate it out. And then we go on diets self and do changes to its specific to ballistics, so the first one is that on day voltage regulation.

23:38

So when you look at X M P, you're hitting this. This memory is 1.3, 5 volts or more or 1.5 volts.

23:46

In the case of our Ballistics Max and we really want this product to utilize all that voltage. when you look at it ... part, you're wanting it to sift voltage, we don't want that, we want this stuff to absolutely use every amount of voltage we can throw at it so we can get the speed and the performance out of it.

24:03

So, we actually go on the day itself, and we unlock that voltage regulator that's built-in, and let it let it basically use all that voltage when when we hit it with that voltage, there's internal timing arrays.

24:15

Our tune as well. So, the best way I can describe this is, if you wanted to get from Point A to Point B, Genetics says you have to go up the road. Hit the stop sign maker, right?

24:26

Hit another stop sign maker, right? to get from Point A to point B to get around the block.

24:30

We don't want to do that. We're not concerned necessarily with Jetpack or how things are laid out. We're actually cutting a path right through the field to get from point A to point B So you have better signal integrity. Your trace links internally are much smaller. And that translates to performance.

24:47

And then the last thing I'll point out here is what we're doing with the PCBs. So with Micron, we can engineer the PCBs specific to the product.

24:56

We can take that DRAM and cram it as close to the Edge connector as possible. So you have your trace links very short, again, signal integrity, point A to point B, and you end up with a very fast part. So as you can see, all this stuff culminates into amazing performance and also consistency.

25:15

So a lot of times we'll see other people in the market saying, you know, you got to buy it as a matched kit. We would never consider having to do that for our server product.

25:25

Nor do we want to do it for gamers. So we're consistent. And we're fast. And we're engineered for gaming.

25:31

And that really is what sets us apart. And we're excited to continue to go after these world records and and provide awesome gaming DRAM for those out there.

25:41

And I believe that's the end of my deck. I do have to leave after this.

25:46

So, if there's any questions upfront, would be happy to answer them. Or if there's something you think of later, by all means, we can make sure that you guys have our e-mails, You can reach out, and I can answer questions. Yeah. So, we do have some, some questions, and, yeah, like Ryan just said, everybody, if you have a question, and, you know, Ryan is not here to be able to answer it. We do gather all the questions and provide them back to microns, so they'll be able to follow up with you through e-mail.

26:16

So, if you have questions and we don't get a chance to ask it, you know, Don't, don't worry. We'll be able to make sure that Micron gets that question over. So, few questions in the queue right now, one, which I was actually wondering about, as well, because Andrew mentioned DDR, or DDR five, and I know that transition is dependent on motherboard.

26:38

And other factors, but in your world, for more Micron sitz, I mean, how do you see that transition taking place as we move, From PDR, Florida DDR five, Do we know what the timeline looks for that and availability and those kinds of things?

26:54

Yes, certainly. Know, right now, we really can't say a whole lot about DDR five. Obviously there's a lot of MBAs in place. You know, I can say that, we definitely want to make sure we're competitive in the market, and we want to make sure that we have the performance out there. So, when this transition does take place, we are absolutely going to be ready for it.

27:14

Alright? So, we know on the server side that ECC is a real important feature inside of servers. There's been more and more discussion about ECC, having a lot of value in the desktop space, and in gaming.

27:30

Do you see more ECC generation and memory coming out for for that piece of the market?

27:38

And you see more motherboards kinda in the client gaming space that will support UCC coming out in the future as well.

27:46

Yeah, so, so that's an interesting question. It's definitely something that we have no way back and forth. It kind of touches to that, to that, that Max product, right, where it's not just for gaming. You look at the content creation and some other applications where you may have a Thread river system that you've built that. You're doing content creation or AutoCAD or something like that.

28:07

And you go, hey, maybe you're a gamer as well, or you just want to look at, no, we don't have it today.

28:13

But, but it's, it's a, it's a great question, it's definitely something that we're seeing more and more interest of a ECC capable product that kind of fits into that performance space. So, a heat spreaders, if you will, ECC gaming product performance from it.

28:29

Yeah, great question.

28:31

We are seeing no more asks for that product. You don't have something today, but who knows what the future holds. So, we all know that for memory. You know, you can over clock memory, and you can push the performance of the memory, But are there advantages to ability to under clock the memory? Like if you under clock The Memory, could you extend the life of the memory and is that even possible to under clock?

29:00

Yeah, it is. You know, I think that a lot of this memory, when you really look at it, it has so much longevity built into it.

29:09

Certainly an over cloche condition could potentially stress components.

29:15

I would say that even even even at a fairly healthy over clock or just an X and P setting, you know, that that product's going to far outweigh the life of your system or even the capability of it.

29:28

When you look at the PC cycle, yeah, that's an interesting question.

29:32

To answer it upfront. Yeah, yeah. I think it would extend the life. Does it make sense to do that? I don't, I don't think so.

29:39

Get that performance realize it, you know, let it work for you, and even at an ..., or as a light over clock, you should have plenty of longevity in the product.

29:51

All right?

29:52

And for those of us on the line who might not be as familiar with PCC, can you touch really quickly on what, what is ECC and what's the benefit of ECC?

30:06

Yeah, actually, and I will punt that one over to Jake because I will I will word salad that one. I know. I know a little bit about it. Jay can get very technical about it and if he's willing to talk to it.

30:20

I'll push that over to M, Right. And that's actually perfect timing, because I think Jake is our next presenter. Right? So why don't we lead into that? Let you answer that question, and then you can, you can kick off your presentation. So does that work?

30:37

Yeah, that works just fine. So thank you so much with the slides as we put them together. So thank you.

30:45

OK, so my name is Jake Lockhart. I'm actually the DRAM product line manager for all crucial DRM. So worked very closely with Ryan.

30:53

We have two other product marketing managers who are not here today. So I'm presenting their product lines.

31:01

So, basically, any any crucial product you see out there, I'm involved at some point, and a lot of my work is actually with Ryan where he says, I want to go this fast, and I say, OK, Give me a couple of minutes. I'll see what I can get.

31:13

So, lots of fun working together, especially on gaming, so let's talk consumer.

31:18

Next slide, please.

31:23

OK, so I'm sure we have all been through many cycles of the experience where you buy a laptop, and you look at your budget, and you decide what hardware you want in there, and how much you can afford, versus how much you would really want, and you come home with that.

31:37

Nice, shiny laptop that has eight gigabytes of memory, and you think there is no way I will ever need more than eight gigabytes, this is more than enough for me.

31:45

And then two years later, when content, all the content on the web has advanced, and there's more content, higher resolution videos, more videos, all of a sudden, your laptop that you thought would never, ever need an upgrade is the slowest machine you've ever used.

32:01

And so that's where a crucial consumer comes into play for, or upgrades.

32:06

It's the fastest, easiest way to increase your performance, because you know, whatever content you were put using your laptop for two years ago that were eight gigabytes was more than enough memory.

32:17

Now all of a sudden you need 16 gigabytes, or maybe even up to 32 gigabytes.

32:22

And, you know, provided you have a system that has available slots, the upgrade is pretty easy.

32:28

Turn off the laptop, put it in the memory, and you're ready to go.

32:32

And so a few things about crucial memory that are very beneficial to that process.

32:37

I mean, number one I already touched on, is easy to install.

32:41

Number two compatibility is Top concern or Crucial Memory.

32:45

What we mean, and there's several areas of compatibility. I mean, number one, does it boot at the speed that it is rated for. That's a very important piece. Nobody wants to buy at 3200 module, take it home, and it runs at 2133.

32:58

Second piece of that is backwards compatibility. one of the things we've seen in the market for this, even even more complicated the past couple of years.

33:07

And you'll see this on the next chart, is the speed that your system is going to run at Can become very convoluted depending on the number of modules you have in your system. That type of processor that you have purchased, the bias that's installed on your system, and the number of ranks per dam.

33:23

There's this all equation that comes into play that says, OK, if you have A, B, C, and D, and you have to down clock one speed, right.

33:30

If you have E on top of that, you have to down clock to speak, right. And so it's not as intuitive as it used to be, or you could just buy a module at a certain speed in your system would run at that speed. It's more complicated now.

33:41

And so one of the things that crucial memory has built-in is guaranteed backwards compatibility.

33:47

There's actually data that we program into a little chip on this module that the system reads that has a full matrix of speeds and latencies that are supported.

33:56

So that regardless of what system you're putting this memory in, whether it's a high-end system that wants to run full speed, a lower end system, or an older system, Our memory will work at whatever speed that memory controller wants to run at.

34:10

And so you get a lot of compatibility with this.

34:12

And we have tools online that actually can go in, scan your system, and we'll talk about them a little bit later and make recommendations for what you should install on that system.

34:23

Now, everything I'm talking about on this slide is non ECC.

34:27

So, Pursuant to the question about ECC, this particular member crucial Desktop and Laptop Memory is all non ECC.

34:34

So no error correction built-in, but it comes in the form factors like ..., for Desktops and ... for laptops.

34:44

All right, next slide.

34:47

So like I said, Market is becoming very cumbersome with so many processors on so many different speeds that are supportive.

34:54

And so what I've put together here is, I've tried to put together a chart that basically has everything you would hear from a customer.

35:02

So if a customer says, I have an 11th gen processor, or I have is then two core, or no matter what, it is, whatever term they're using.

35:09

Whether it's the name of the core With A name of a, code name of the process, or the marketing name, of the processor, whatever, it's all listed here.

35:17

And then I tie every one of these to a memory, A memory speed, which is, again, it's not guaranteed at that speed. It depends on the Dens per Channel and the ranks for them, all of those equations. But in general, what is the memory attached to all these processors?

35:33

And then what do we recommend for for crucial? So just to start in the middle, we've got Core eighth and core ninth gen, referred to as coffee, like I'm talking like refresh.

35:42

Got rise and Gen one, referred to as Summit Ridge and Raven Ridge or Zen Architecture.

35:49

All of those processors, all those names are associated with 2666 memory.

35:54

And crucial has a full portfolio of 2666 products.

35:58

We also have our Ballistics products at 3030, 200, up to 4000, that are really good candidates for these types of processors, meaning the memory controllers are, in general, strong enough to be able to hit 3030, 200 and up to 44,000 if you have a high-end system.

36:16

And so a similar approach here, even going down to core 11th gen, Intel's rocket lake that, hopefully, you heard about this week An AMD Rising Gen three, Gen four. And Jen, five, all of those processes are associated with 30, 200 memory.

36:30

That is fully backwards compatible with every speed and DVR for.

36:33

So, if you want, if your customers want the future proof, regardless of their systems, by our 30 to 100 memory, and it will support any speed, that the, the older systems want to use, it, doesn't matter.

36:46

So, I know there's a lot of data on this, but, unfortunately, there's a lot of terms in the industry that become very confusing.

36:53

And soda.

36:55

We wanted to provide, this as a reference, so you have a way to correlate all of those terms to a specific kind of memory.

37:02

So I'm, I'm gonna pause here, this is my last slide for consumer, any questions on consumer DRAM.

37:10

So let me take a quick look, but you know, I'd do have a quick question.

37:18

You were talking about the guaranteed backward compatibility, which is, which is a really cool feature, but I got a little confused, is that, only on the DDR for 3200 or all.

37:31

Crucial memory is backward, compatible, and guaranteed.

37:36

All crucial memory is backwards compatible to speeds lower than the rated speed. So, if you buy a 2666 part number, it is backwards compatible to 4533.

37:47

Alright, so I talked about that, that matrix, that of speeds and latencies. And that's what we do every one of these products.

37:56

It includes a matrix of everything that's supported.

38:00

That is backwards compatible.

38:02

We've done competitive analysis and one of the things we've found from our competitors, if you, you go buy a 26, 66 memory from a competitor.

38:10

There's one value in that matrix is, of course, 2666. That's it.

38:14

So if you get a memory controller that doesn't like that speed, you can run into compatibility issues with our memory.

38:20

We program a full, every possible combination into that matrix so that you, no matter what system, you install it and it's able to run.

38:31

OK, so, I've got a couple other questions here. I'm, I'm, I'm hesitating because I need to make sure that I read them correctly. Otherwise, I don't know if, they're gonna make sense though, or the consumer memory.

38:48

What's the max amount of memory that can be installed before you have to resort to registered?

38:54

And I assume that means before you have to resort to registered memory, correct? Who registered?

39:00

So, yeah, you're reading that correctly.

39:03

And MDD or for for consumer applications, the maximum density per module is 32 gigabyte and that comes into the and that requires a 16 gigabit die.

39:15

And so it's a to rank by module with a total of 32 gigabytes.

39:21

Now, it's It's kind of It can get a little bit cumbersome, because it depends on the base component you're using.

39:28

So if you see up here at the top of the chart, Rachel, if you can circle in the black bar where it has eight gigabit and then the three bonzo configurations underneath it.

39:40

Go up towards the top, to the right. David, you're muted if you're trying to say something.

39:48

OK, yeah, David's driving the slides. Oh, sorry, David.

39:52

So over towards the right-hand side, where it has in the header bar, eight gigabit, and then it's four gigabyte, eight gigabyte, 16 gigabyte.

40:00

So it depends on the density of the DRAM components. So if you have eight gigabit components, the maximum density per module is 16 gigabyte.

40:10

If you have 16 gigabit components, the maximum density is 32 gigabytes.

40:13

So that's why I say, you know, to ask the question directly, it's 32 gigabytes.

40:20

It does depend on the component a year that is being placed on the module, so it's it's a little bit it can get convoluted there, OK?

40:29

All right. I think we can go ahead and keep moving forward. I know there'll be another opportunity here at the end of your presentation for more questions. So go ahead and keep sending in the questions.

40:40

We'll let Jake continue with his presentation, and we'll pick them up at the at the end.

40:47

Thank you. Next slide.

40:50

So we're going to dive into server DRAM So you know, before I go through this slide, I didn't want to touch on the question of, you know, what error correction, what is it, what do we mean, and what are the benefits?

41:00

So, basically, what we're talking about when we refer to DRAM memory correction is we're actually adding a addition one, an additional component per rank.

41:12

So, if you have a single rank module, it has 1 or 2 components depending on the configuration.

41:18

If you have a dual rank module, it's either 2 or 4 extra components.

41:22

And those components are included on the module so that the memory controllers can perform error correction algorithms to verify that the integrity of the data, um, there's mechanisms that can affect DRAM usually in the end, in the public interest either referred to as alpha particles.

41:43

Gamma rays is a cosmic type of events that can at times cause DRAM cells to want to flip their state, and which has the potential to corrupt data.

41:55

And so these extra components are put on there, so that the member controllers can calculate, to make sure is the data correct and if not, correct it.

42:05

And so, when we talk about ..., which is already brought up, we've got All of those have extra components on the module to allow that to happen.

42:17

Now, a couple of things are required there. Number one, you obviously have to have a module with the extra components.

42:23

Number two, you have to have a system and a memory controller and a ...

42:27

that understands error correction and is programmed to do that.

42:31

And, the reason I bring that up is that there was a great question asked to Ryan about ECC, entering the gaming market, and he's absolutely they, right?

42:39

We have been looking at it, it largely depends on, do the processors support it and do the motherboard vendors integrate it into the vials?

42:47

Because if you take an ECU dim and you plug it in a system that does not know what ACC is, it ignores that extra component and just runs it as a non ECC.

42:56

That's obviously not what you want.

42:58

And so there's a large dependence there on the bios and the CPU features.

43:03

And I would say for the for the desktop market and I'm talking about standard desktop, gaming desktops, content creation, anything that's a desktop type platform, we're still kind of in the very early stages of that being integrated as a feature that's available that could result in a gaming ECC type product for servers ECC has been around for a long time.

43:27

And that's always integrated data server.

43:29

So when you talk about CPUs like Intel Xeon, AMD epic, those are built.

43:35

Those always have error correction because that's critical to a server to make sure it can stay running 24, 7 without have any errors that cause the server to have to go down.

43:49

So I'm going to pause there. Did that answer the question about error correction?

43:55

Yeah, I think that was, that was perfect. So thank you.

44:00

So talking about server memory. We have lots of different forms of this memory available.

44:06

You know up to and including even VLP are damaged VLP ECC UTM VLP stands for very Low Profile.

44:14

So, if you imagine a standard height module about 30 mm high VLT, these are about 18 to 20 mm, so therefore, the low low profile, 1 U 2, 1.5, you blade servers that just need everything as, as low profile as possible, and so, we do have those available.

44:33

This memory benefits from the same vertical integration with micron, meaning that we are designed with backwards compatibility in mind.

44:41

And so, these, just like our consumer products, our server products are 100% backwards compatible to slower speeds.

44:50

They'll also run at full speed and the goals be they're rated for without any issues.

44:56

I think we just lost the presentation.

44:59

This is not coming up on my screen, insurance broker.

45:14

There we go.

45:14

Yep, that's better.

45:16

OK, so, like I mentioned, all server modules have ECC. They're all built with that in mind and the processors know what to do with them.

45:24

We do a lot of quality and reliability testing with these components because we know where they're going into an application that needs to be reliable long term.

45:33

These data centers have racks and racks of servers and storage, or, you know, all sorts of data from all over the world, From social, social media, all the way up to Enterprise banking.

45:45

And so these servers have got to be reliable, and we do a lot of reliability testing to make sure that our parts are able to support that record environment.

45:55

We also support the Mac pro-line with this same high quality.

45:59

And so, as you've seen the new macro systems launch, those products benefit from the same micro.

46:07

On.

46:10

Liability standards.

46:15

Let's jump to the next slide.

46:24

There we go. So, one more memory to speed to CPU reference, Samus consumer here, I'm trying to throw out every name you might hear as you talk to customers about what type of server they are, installing memory and whether it's an older server or a brand new server.

46:39

And so, you know, up through 29 33, 3400 with Intel, Cooper Lake, and AMD Realms, respectively, we, there are new generation systems coming online that are expected to launch this year. I call them until Next Gen at AMD Next Gen, because those names aren't quite yet public, but we'll update this chart as those names go public.

46:59

And the biggest thing to call out here is that the, no, to the question earlier about Max Densities, something to keep in mind here that the 16 Gigabit based products, for some of the earlier systems, are not supported. And so, that max density per slot could mean something different, depending on the type of hardware you have.

47:20

So, it's something to keep in mind as you're talking to customers, about where their memory is going, to make sure that the correct compatible product is installed.

47:30

We do support all configurations of ardan's ..., you can see them listed at the top, black bar.

47:39

Because, you know, we don't want to sell a thousand part numbers. But, to be honest, every one of these configurations has a place in the market.

47:47

Some are higher volume than others, but each of them has a specific reason for why it's configured the way it is.

47:53

So if you need any help with one of your customers on a recommendation for either for our configuration module, let's say they come and say, I just want a 32 gigabyte module. Tell me what to get, help us, understand what their application is, and what hardware they have, and we can make them more detailed recommendations for them.

48:12

We're happy to help there.

48:16

All right, and that's everything I put together for Server. So I'll open it up for questions. OK, great. I think one of the first questions that we had early on is we want to know where you got your goodies.

48:30

Where do you get that all?

48:33

Yeah, I'll be honest. I forgot, I was going to be on camera today, but I'm glad I want to order. these, internally.

48:38

It's actually, a couple of years ago, we had a product that was co branded with ACS, the aces Tough Line and we designed the heat spreader that kinda fit the same theme.

48:50

We had, uh, you know, jerseys and hoodies that we ordered all with the same ... of black and gold.

48:56

So I also call this my World Record Hoodie. I'm like Ryan said, we broken five world records with Ballistics and sell them.

49:04

Is my gold medal Jerzy?

49:05

I wear this as often as I can.

49:10

So a couple of questions coming in are not related to the hoodie, but thanks for letting us know.

49:19

What can you describe, what's the difference between an ...?

49:25

So, you dhamma UTM is short for un buffer.

49:31

And what an buffered dim means is there are no buffers between the Memory controller and the DM components?

49:37

There's no extra components that are trying to improve signal integrity that are performing any operations. It's a direct connection between the Memory Controller and DRAM components.

49:48

In the case of an R dam or otherwise known as a registered dam, there is a component called a register or an RCT that actually all of the command and address signal's route through the register first and then it distributes to the DRAM from there.

50:04

And you might be saying, well the UTM sounds great.

50:07

I don't want any components, to, you know, adds latency or caused delays, and the register actually doesn't cause any problems along those lines.

50:15

It's there to basically manage the command and address traffic for higher density dams.

50:22

And so in the case of ..., the maximum number of components that you can have on a UTM is 16, In the case of art hymns, you can have up to 36.

50:31

And we even when you start stacking the around components, you can have up to 144, that's a lot of components for a memory controller to manage.

50:39

And so when we get into those higher, higher density modules, that's where the register really comes in to play, is that it's there to directly control the command and address signals to keep all the and to keep all the DRAM in sync.

50:53

In addition, the register also performs error correction: operations on the command and address and so if you can imagine in a server application, obviously, you don't want your data to be corrupt But you also don't want your command or your address signal to get corrupt either, and so they're registered provides error correction for the grant an address And then the DRAM has set up in the servers to Provide the data correction.

51:16

So it's really a combination of more DRAM, Meet needs, another component to keep everything in sync.

51:23

and we need to have error correction on the command that address and high reliability patients.

51:29

OK, so I'm kind of looking at the charts even showing and putting up with your, with the compatibility between the different processor families.

51:39

We can see that the memory for the server looks like the largest capacity are shown as 128 gigabyte.

51:46

Are there Larger memory modules available, do we have 256 gigabyte or removing to 512 what's kind of the future look like for no capacity for the server memory?

51:59

That's a great question. So, yes, there is a 256 gigabyte supported option, in DDR four.

52:07

That's actually good feedback.

52:08

I might need for this crucial and Micron actually do not offer that product. It's such a niche market at such low volume and the cost is so high to stack the DRAM from that. The Micron made the decision to not support that product.

52:23

Having pursuant to your question Yes. In DVR for with the 16 gigabit component, you can achieve up to 256 gigabytes per module.

52:33

That's the max for DDR for as we look ahead to DDR five, obviously see that higher density components and so you will see higher densities, eventually in DDR 5 beyond 256 gigabytes.

52:48

Probably a lot at launch, it will be 256 gigabyte max, all right. So in looking at, you know, Intel's isolate processor, and I'm hoping you can answer this question, I know we might be, you know, a little bit ahead with ice like but looking at the isolate, are you getting better performance on to our dims over one arc? And if there is, you know what the differences?

53:14

Oh, that's a good question.

53:15

I haven't seen any data specifically from ice like in that way, but I can answer the question generically that says it depends on the workload at all depends on what the server is doing, and the reason I answer that way is.

53:28

But in servers, you know, number of dams, it's actually not quite as important as the number of ranks of memory.

53:36

And so, you know, a lot of these server dams are dual rank, meaning they have two groups of DRAM on there.

53:44

And if you have a workload that can leverage, that can take advantage of those multiple ranks, meaning it's performing different operations on different ranks at the same time, it can really benefit from more ranks. So, you know, in your example one dam versus two dams.

54:00

Now there's a lot of cases where two teams will do better, simply because the memory controller has more ranks to play around with. Therefore, it has more options to do work in parallel.

54:09

Now, if you have a workload that doesn't understand, the number cannot leverage all of the ranks and your server. And in that scenario, one rank versus two ranks might look identical, because it

doesn't know what are sorry, one Denver says to them. I look the same because it doesn't know what to do with the extra ranks anyway.

54:26

So it's very workload dependent.

54:30

And that probably doesn't answer the question in the way you want it, But unfortunately, that's the way it is right now, And cloud computing.

54:38

Yeah. That was a great answer. So, before I kind of keep going, I just want to make sure, are we at the end of the micron deck, is that, or there are more presentations to go from here.

54:53

I think the only one after this is on our compatibility tools, and it's a It's a pretty brief. Yeah, the Micron. Selectors.

55:00

OK, so why don't we go ahead and do that? We're coming up near the top of the hour, so let's go ahead and finish that out, and then we'll close with a few additional questions, and then we can, we can kinda go. I don't know if Andrew is able to come back in or not.

55:18

We might have a couple of SSP questions in here, we want to hit him up with.

55:22

But let's go ahead and finish, and everybody online keeps sending in your questions.

55:27

After Jake is done, we'll finish up with a couple final questions, and then we'll do our wrap for the day.

55:33

So, OK, Sounds good. And I'll ask you to jump in here on the tools as well.

55:40

I mean, these tools are for memory and for DRAM and for SSP.

55:45

And basically, what we're trying to do is provide you with any asset that you would need to figure out what to buy.

55:52

I mean, you can probably already tell from those charts I put together of all the different processor names and all the different memory.

55:57

There's no way you're going to memorize that.

55:59

You know, keep a copy with it, but it's a lot of information.

56:02

Same thing on SSD is there's a lot of densities, a lot of interfaces is a lot of information.

56:06

And so, we, we've provided all of these tools. Whether you want to have your system scanned or you want to put the information in and have it tell you what is compatible, or you think you know what part number you're looking for. And you just need to find the actual one, or you need to configure a server, whatever it is. We've got tools available that will tell you.

56:26

This is the part number, this is the product, This is what you need to install.

56:29

And so, I would encourage you to take advantage of those.

56:32

Whether four: with your customers, your home systems, whatever.

56:37

We see a lot of you a lot of traffic on these and it's all of these tools are our world-class tools to provide you with information you need on what product to purchase.

56:48

So, that's that's my my spiel Andrew, anything to add from the FFT side.

56:54

Nope. I think you got it.

56:57

Alright, great, Let's go ahead and I'm gonna give you, we got a couple more questions in here that I do want to make sure that we get asked.

57:05

So I'm gonna go ahead and field a couple more questions here, and then we'll wrap everything up and give everybody their day back.

57:13

But on the guaranteed backward compatibility, I kinda want to jump back to that just for a minute.

57:22

Is that only on micron memory, or can you mix cease with brands, like, for example, could you mix it with Kingston memory, and in your memory would still be backward compatible?

57:34

You never want to mix memory. You only want to use crew there is the right answer.

57:42

You know, so reality is in DRAM wear and tear on world, and with the Jedi consortium, now these spectra standardized.

57:49

And a lot of cases, now, it's up to each memory vendor to determine what specs are going to support and so, you know, with micron being vertically, or with crucial being vertically integrated with Micron. We know by design what the components are capable of it.

58:02

So, we're able to reliably program in that we can support NSP backwards compatible.

58:08

Now, our competitors, some of them don't do that, because they they, they just buy up a bunch of DRAM build modules and a program with them for one speed others.

58:16

You know, Kingston's a little more innovative than, that, you know, there, they can go test and verify and decide what they're going to program into into the module, and, so, I would say, every memory vendor, everyone that builds modules out there is a little bit different.

58:30

Um, but, you know, in general, when you combine memory modules, the least common denominator wins.

58:36

And so if you mix a crucial module with another module that only supports 2133, the memory controller can't run at two different speeds. It has to take the lowest common denominator.

58:49

And so in that case, if the non crucial module is a 21, 33 only the memory controller will say, Sorry. There's two modules here, but only one of them can only to support 2133.

59:00

Second module. Can you support that?

59:02

Yes, in the case of crucial, we would say, yes, we can go for it, and everything with down, clock 2133.

59:07

So, and you'll get compatibility.

59:09

Just keep in mind that if you mix vendors that don't have exactly the same advanced matrix that we do of speed and latency is you could suffer additional performance penalties.

59:21

Alright, so, I'm gonna give Jacob break here for, for a second, and catch your breath, with the mask on and everything. And I'm gonna ask Andrew a couple of questions.

59:33

So, Andrew, on the data 5300 SSD, what's the expected life on that? And then, on the new NVME drive, that would be replacing it.

59:50

You have to unmute yourself, a good question, and it's something that, you know, we can talk to you more in depth about, in kind of a broad audience here. I don't wanna go too much into our roadmap. But what I can tell you is that the 5300 will have full support through this year, and the vast majority of next year. And we are, we're investigating, And we can tell you more offline, on, Assad, subsequent ... products on the satellite. We expect to be able to support our customers multiple years going forward. And then, on the NVME side, if you're, if you're you're a customer or wanting to move off of sada India NVME, it's our 7300. Those are going to be very, very similar in form, fit function, and they're going to be very close in price. There'll be a little bit of a premium on the 7300, but it's not. it's not huge.

1:00:43

OK.

1:00:43

So on the endurance rating, for the drives, you know, there's several different ratings that kinda get used by different different companies for showing the endurance of their drive, but is there a good place to go?

1:00:58

Look for micron to find things like T DW or, do you have pages show comparison between your endurance and competitors endurance. Especially if you guys are using different ratings. I mean, how would you compare those?

1:01:14

So, I will say, so. So, the key difference in ratings is going to be between consumer and client and data center.

1:01:22

So, those are just specified, even though the metric comes out sounding the same, and drive writes per day, or T B W. Is the workload used to derive that metric is very different. So so first thing is never compare those two. They're not comparable.

1:01:39

Within the data center market and within the client to consumer market, typically, they're going to be very close in terms of the world. There's actually a ...

1:01:49

standard JST, too, 2, 0, 1, 8 0 to 18. I mean, that specifies that and we typically all follow it. And so if you look at a Micron data center SSD and want to compare that to a Samsung, whether it's in TV W or rice per day, I would I would save them. For the most part, they're apples to apples.

1:02:08

Every once in a while, you'll have somebody like Kingston that will try and specify a data center drive over a three year span, which artificially inflates or dry writes per day. Typically, if you go by T V W, you don't have to worry about things like the warranty. It's just, how much data can you write? Or you can, you can do it all in six months, or you can take five years to do it, right?

1:02:30

So on that side, where you can find it, product briefs, data, data, sheets, really most of our ministerial that, let's say, any sort of specifications and including going to our website. You can go and get that. Typically, we do want to talk about more on TV W because it takes a few of those variables out of the equation, but that's that's how I would look at it. And then, on client and consumer, know it's it's very much a less stringent specification because to be honest, client and consumer workloads, like what I'm doing right now on my laptop, my SSD is sitting there idle. It's not really doing anything.

1:03:06

Where if this was a server, there's a good chance on a server that there will be a lot of activity. That would be going 27. not periodic bursts of data.

1:03:18

OK, so let me ask one more question, Andrew.

1:03:22

And then I kind of have a general one that I'm not sure who's going to want to field regarding warranty but on 2.5 inch SATA drives are their larger capacities. And seven terabyte are we going to see like a 7.6 a terabyte of data?

1:03:40

Yeah, so, so we have 716, both in our 5310. And that, at this point, on the side of market, is a fairly common. As the top class, we were not looking into, and I don't believe our competitors are going much higher, or any higher than that. Simply because we always want to look at is a ratio between how fast you can fill a drive and how big the drivers and when one of the when the assay gets way way higher than then how you can write to it. We don't, we don't want to have a solution where would take weeks and weeks and weeks to be able to fill the drive or if you had in some sort of re configuration being able to recover. So, personally, I don't think you're gonna see SATA really go any higher than that 7.6 A. I could be wrong. But definitely, on the NVME side, the NVME side, same 2.5 inch form factor.

1:04:31

You're gonna see 16 and 30 to my class of assistance, right? So, Jake, we're still getting lots of questions about your hoodies, so you might want to put that up on e-bay.

1:04:43

Maybe people wanted to know what size you could come in and Kenya get it and balu maybe Maybe there's a market, I think that's the next sales incentive, right. I think so that might be your inner via module and good intro for a hoodie.

1:05:04

Might have something going on there. So, last question, Jake. Kinda on warranty for the memory. Maybe you can talk to us really quick about what, what's the warranty from micron crucial on your memory.

1:05:20

That's a good question.

1:05:21

So, on, on, crucial, we offer a lifetime warranty for our consumer products except in regions where lifetime warranties aren't allowed.

1:05:30

And I believe our time there is 10 years from the date of purchase, which is, you know, more than the life cycle of most of these consumer systems.

1:05:38

For server are my R R. Warranty aligns with Micron.

1:05:42

And so, for a server, we actually sell micron branded modules.

1:05:45

Just like on SSD side, it's micron brand that SSDS and the the warranty that micron offers on those is a three year warranty.

1:05:53

Great. Alright, well, I'm gonna go ahead and kind of wrap this up, but I want to ask everybody, just hang on with us for a couple more minutes while we, while we kinda close this out. Just as a reminder for today, For everybody for attending, for attending the entire session.

1:06:10

You're gonna get entered into a raffle drawing for a chance to win either an Apple or Samsung smartwatch, whichever one you prefer and we'll be giving that away and we'll announce the winner of that tomorrow.

1:06:23

Tomorrow we're going to conclude our ASI tech summit with our final presentation, which will actually be done by ASI, and we'll be talking about IOT, industrial computing, and some of the solutions that ASI has that fit into this space.

1:06:41

So a lot of the vendors that we partner with have products that are in this category, I mean, we want to make sure that you guys are all aware of what those are. So, you know, those are things that go into digital signage. There, go into robotics.

1:06:54

They go into healthcare five G, security surveillance, so lots of stuff there that I know you guys are involved in, so if you're thinking five G is not me. Actually, it is.

1:07:06

So, I highly recommend that you all attend tomorrow and also that allows you to get entered into the raffle drawing for our grand prize, which is a all on one computer system.

1:07:16

So, with that, I've done my spiel. I wanted to, you know, thank Team Micron for an awesome presentation. Paul, is there anything you kind of want to say, in closing to our audience of resellers, here.

1:07:31

We wrap. We just appreciate, appreciate all the support and the business, and, again, we're here as a resource, so if you ever need anything, please don't hesitate to reach out.

1:07:39

Obviously, partner strong with your team, We can help you through it.

1:07:44

Great.

1:07:45

All right. Well, again, on behalf of ASI, I want to thank everybody For joining us. It's great to have such a big group and we really appreciate you bringing all those resources together on this event to talk to our customers really, very beneficial and extremely helpful. So on behalf of everybody on the call, thank you to all the resellers and customers and attendees who joined us today. Thank you guys, again, for joining our Day three.

1:08:14

We look forward to seeing you here again, tomorrow.

1:08:16

We will make the presentation, a slide deck available to all you guys and along with the recording.

1:08:23

And with that, I'm going to go ahead and end. It will give you your day back, and we'll see everybody back here tomorrow.

1:08:28

So, thanks again, everyone.